Unshackling Innovation Through Multicore Architectures

UNBOUND COMPUTE™

Scott Sellers

VP Hardware Engineering, CTO, and Co-founder

sellers@azulsystems.com

AZUL SYSTEMS

Today's Agenda

www.azulsystems.com

- What's Driving the Trend Towards Multicore?
- Multicore Landscape
- Multicore Challenges
- Multicore Incarnate: Network Attached Processing and Azul Vega™ Processor

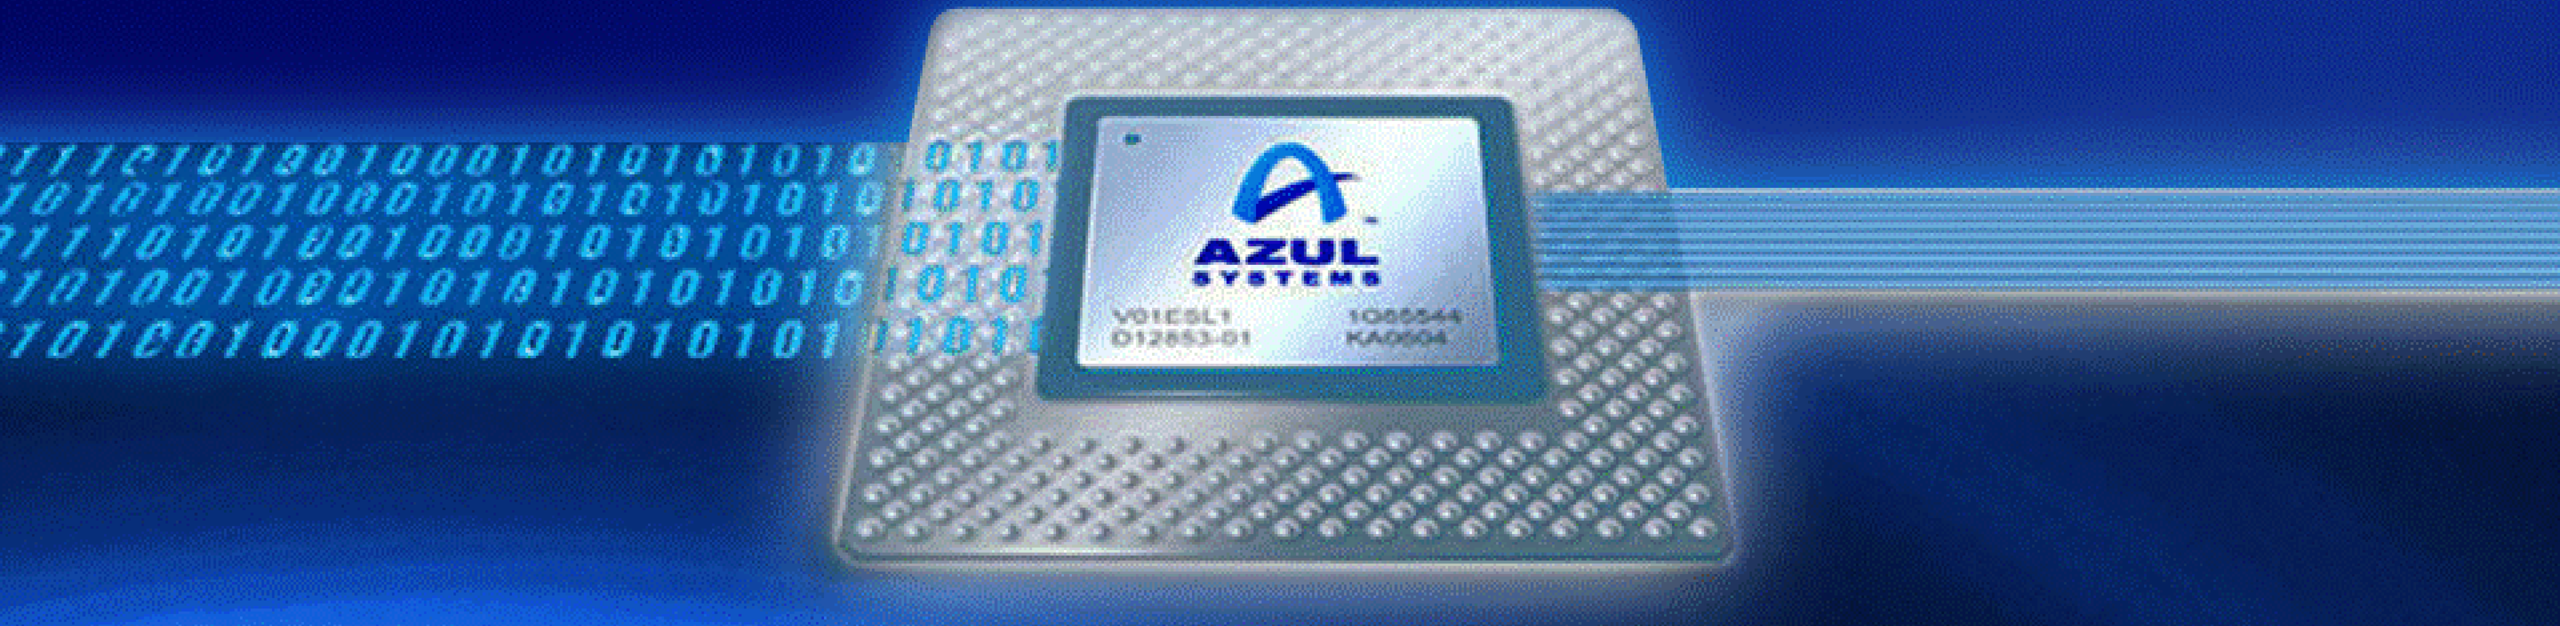

The Thermal Problem

Intel x86 Microprocessor Frequency

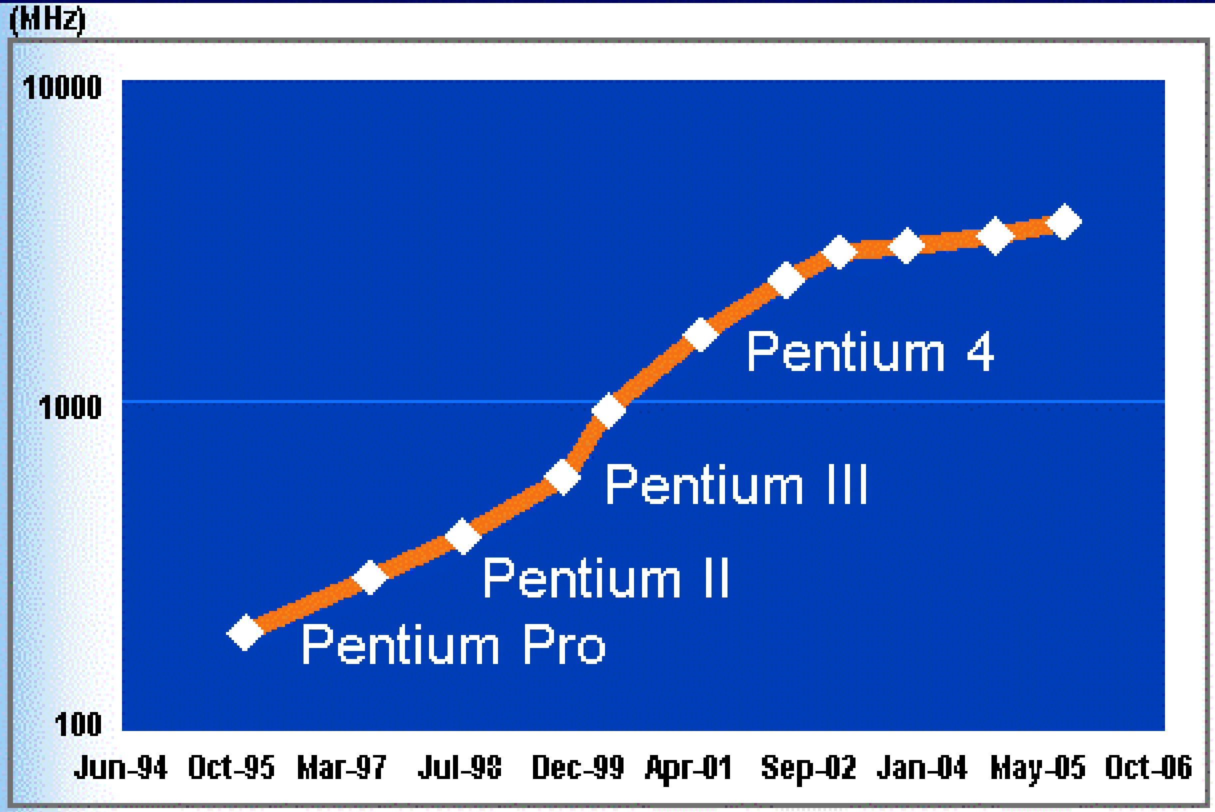

Intel x86 Microprocessor Power

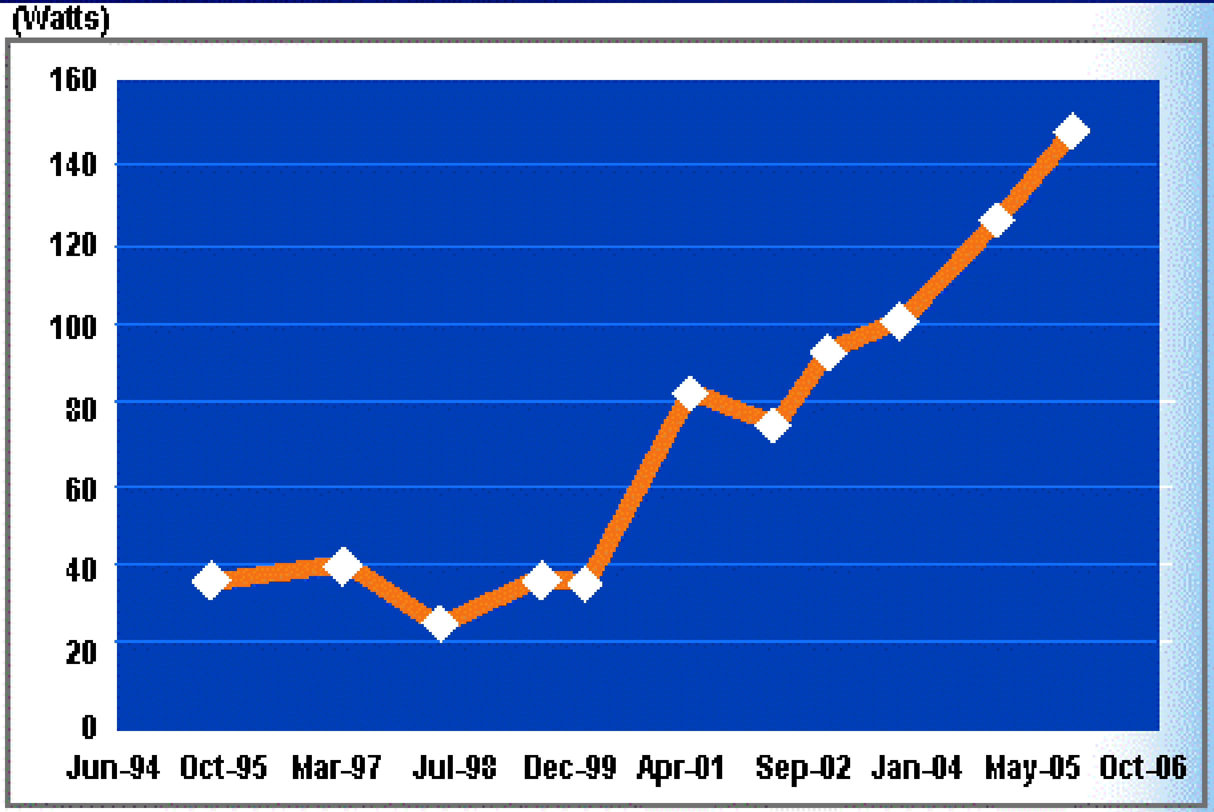

- Since 2002, x86 microprocessor operating frequencies have increased at ~12% annualized rate, compared to the ~40% historical rate
- Power dissipation is clearly limiting performance

The Power Density Problem

AZUL SYSTEMS

www.azulsystems.com

Intel x86 Power Density

Genossar and Shamir, 2003, "Intel Pentium M Processor Power Estimation, Budgeting, Optimization, and Validation", Intel Technology Journal, Vol 7(2)

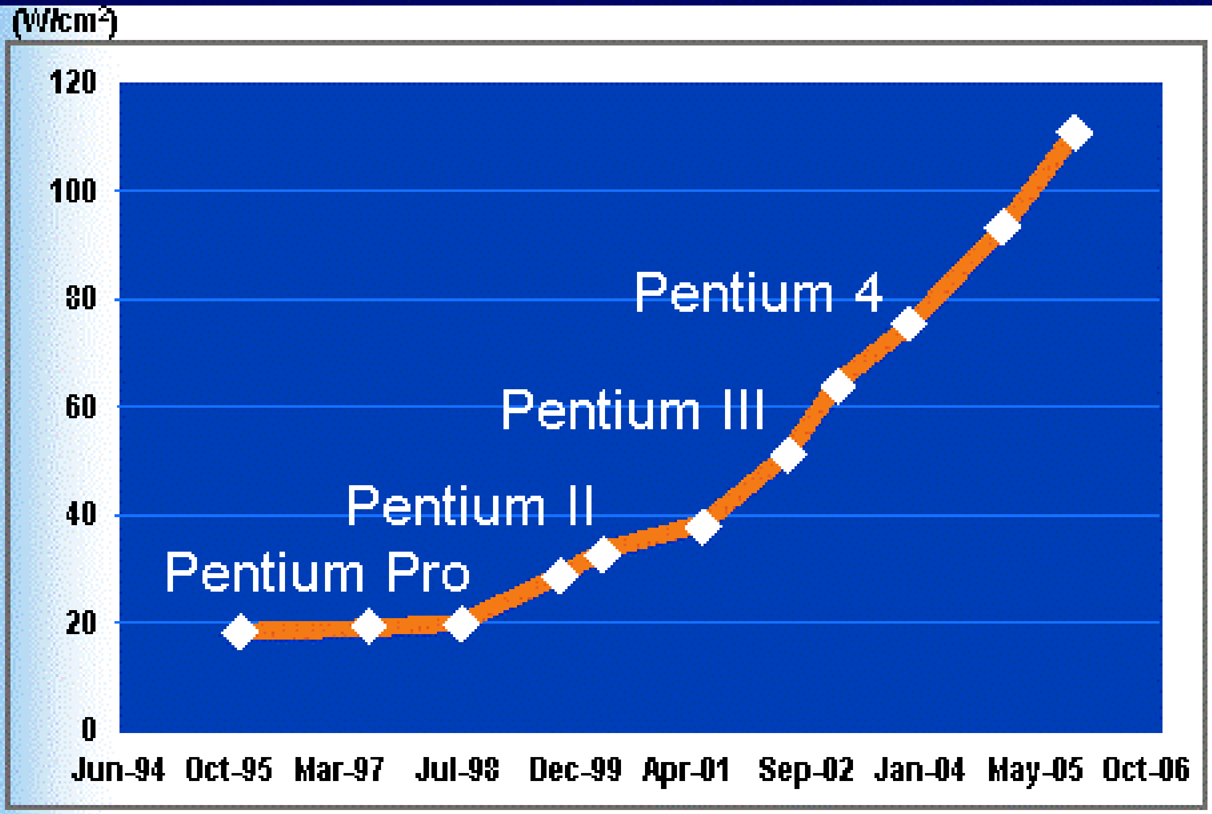

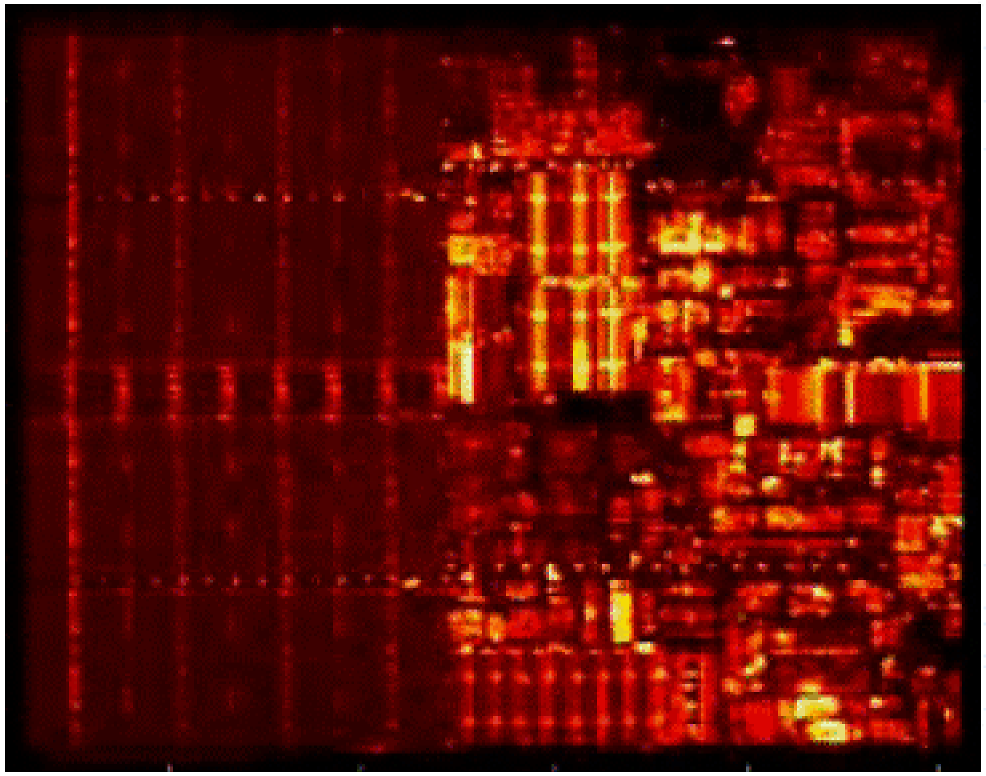

- Dramatic growth in average and peak power densities
- >100 W/cm2 average, ~300-400 W/cm^2 peak
- Air cooling technologies not keeping up with thermal management requirements

Multicore Solution Trends

AZUL SYSTEMS

www.azulsystems.com

Today's Challenges

- Frequency and performance growth has plateaued due to power limitations
- Power densities are at the limit of air cooling
- Gap between CPU and DRAM speeds is limiting performance
- Design time, complexity, and costs are increasing

Multicore Solutions

- Larger number of lower frequency cores
- Optimize for Throughput
- Power-efficient and higher performance function-specific accelerators
- More transactions outstanding (multicore / multi-thread)
- Simpler processor cores
- Smaller design teams and easier validation

PARADIGM SHIFT

Today's Multicore Landscape

	Workload Type	# Cores / Threads per Chip	Max Coherent Threads per System	Scalable Software Development Complexity
IBM Cell	Gaming / Media	9 / 9	1	High (no coherency)
AMD Dual-Core Opteron	General Purpose	2 / 2	16	Medium (ccNUMA)
Broadcom BCM14xx	Embedded / Networking	4 / 4	16	Medium (ccNUMA)
Intel Dual-Core Xeon	General Purpose	2 / 4	16	Low
RMI XLR700	Embedded / Networking	8 / 32	32	Medium (thread scheduling)
Sun Niagara	Enterprise Web & Application Tier	8 / 32	32	Medium (thread scheduling)
IBM POWER5	General Purpose	2 / 4	128	Low
Sun US-IV+	General Purpose	2 / 2	144	Medium (ccNUMA)
Intel IA64 Montecito	General Purpose	2 / 4	256	Medium (ccNUMA)
Azul Vega 1	Enterprise Application Tier	24 / 24	384	Low

Multicore Challenges

Memory Bandwidth

Platform Scaling

Amdahl's Law

Power

Software Licensing

Slower for some apps?

Automatic Parallelization?

Software Development Effort

Where are the multi-threaded Apps?

System and IO Bandwidth

Heterogenous Multicore?

Multicore Power Challenges

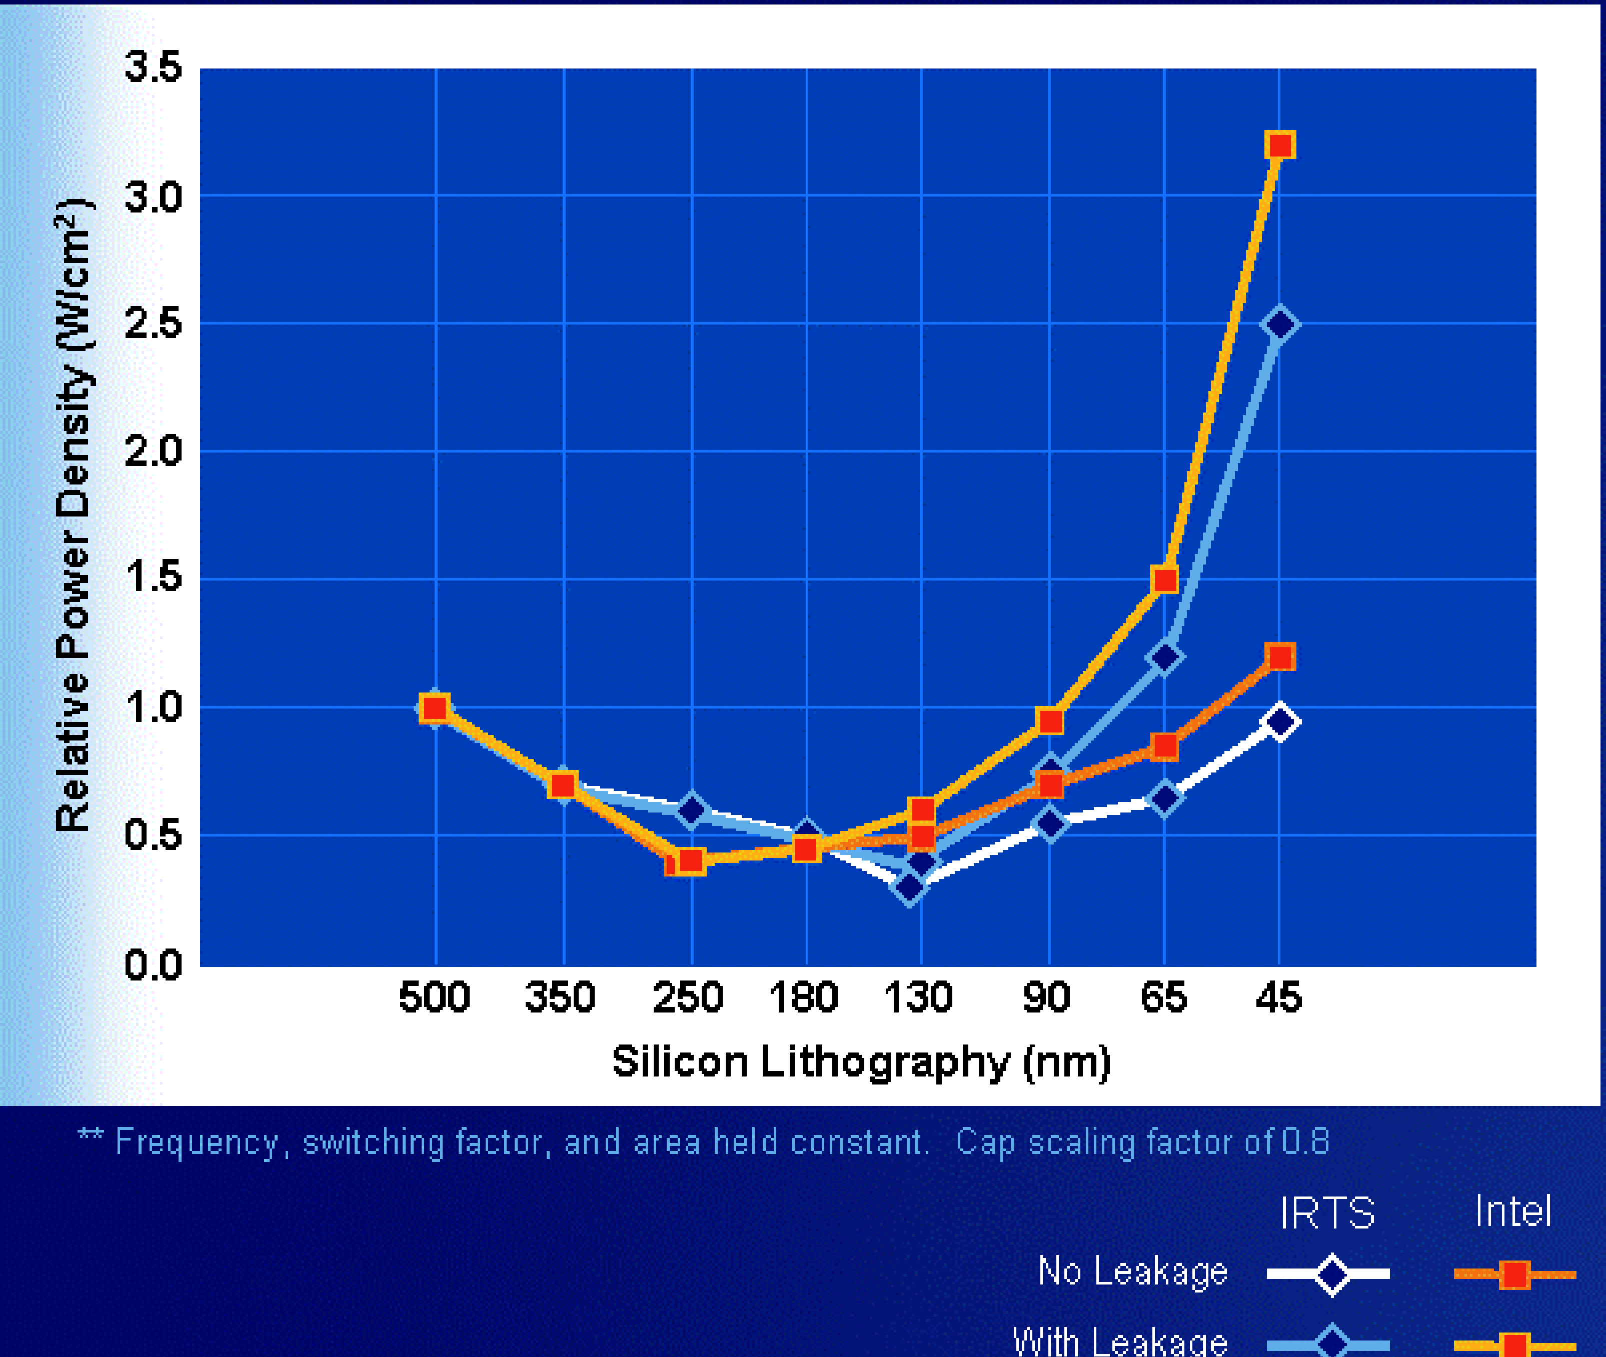

** Frequency, switching factor, and area held constant. Cap scaling factor of 0.8

- Voltage (squared) is no longer scaling at the same rate as density
- Simply adding more cores as lithography allows quickly runs into the same thermal challenges we are already wrestling with today

	Core Voltage	
	IRTS	Intel
500 nm	5.00	5.00
350 nm	3.30	3.30
250 nm	2.50	2.00
180 nm	1.80	1.67
130 nm	1.20	1.44
90 nm	1.15	1.29
65 nm	1.03	1.16
45 nm	0.97	1.08

Amdahl's Law
Serialized Portions of Programs Limit Scale

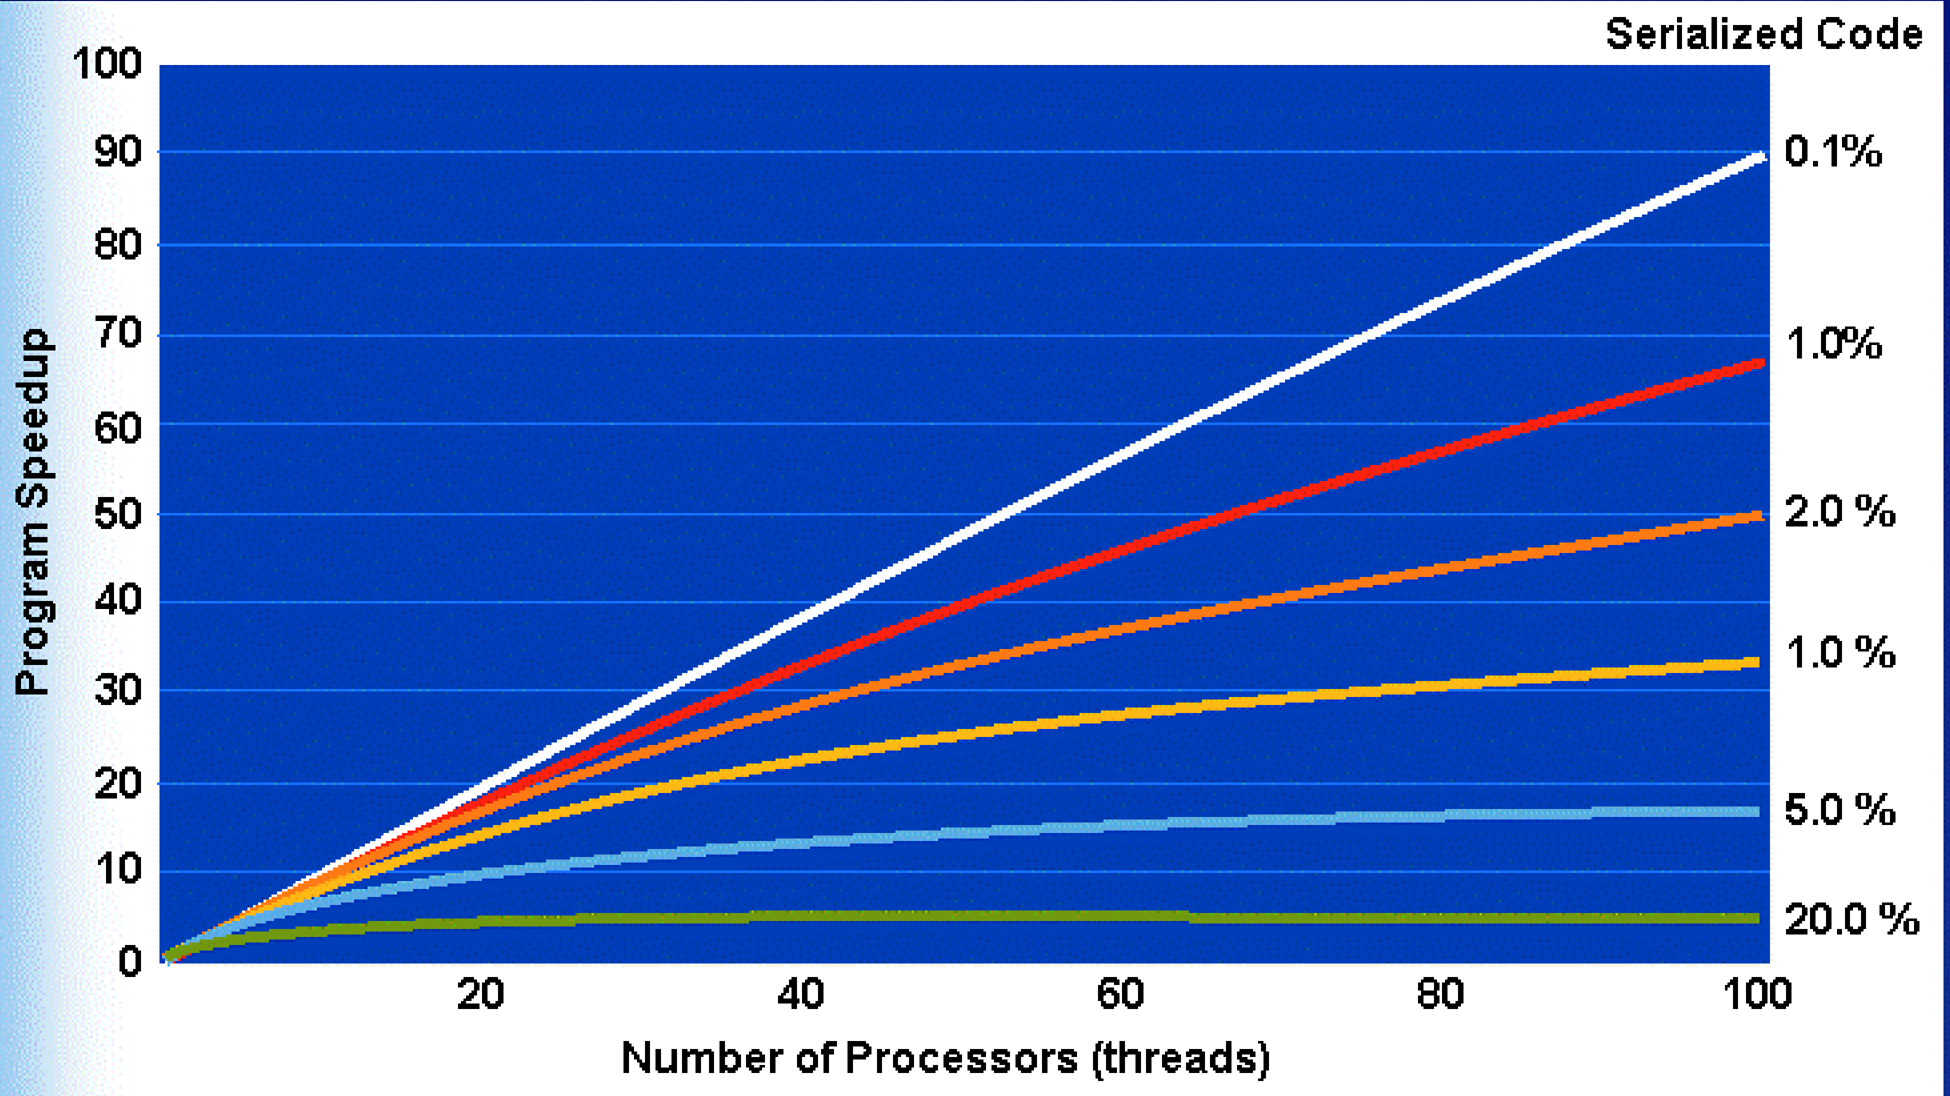

- In 1967, Gene Amdahl discussed "the continued validity of the single processor approach and of the weaknesses of the multiple processor approach"
- Performance benefits of multicore are limited by the number of operations that must be performed sequentially

Multicore Incarnate:
Networked Attached Processing and Azul Vega™ Processor

www.azulsystems.com

AZUL SYSTEMS

Taking A Page From History

Processing | Appliances

Storage | Filer

Routing | Router

1990's

2000's

1980's

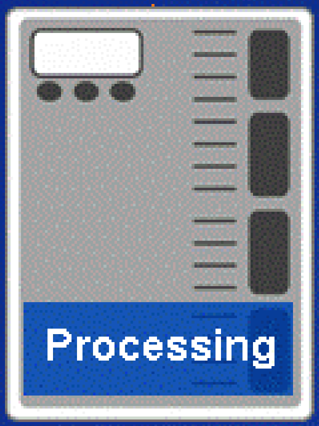

Analogy: Storage Consolidation

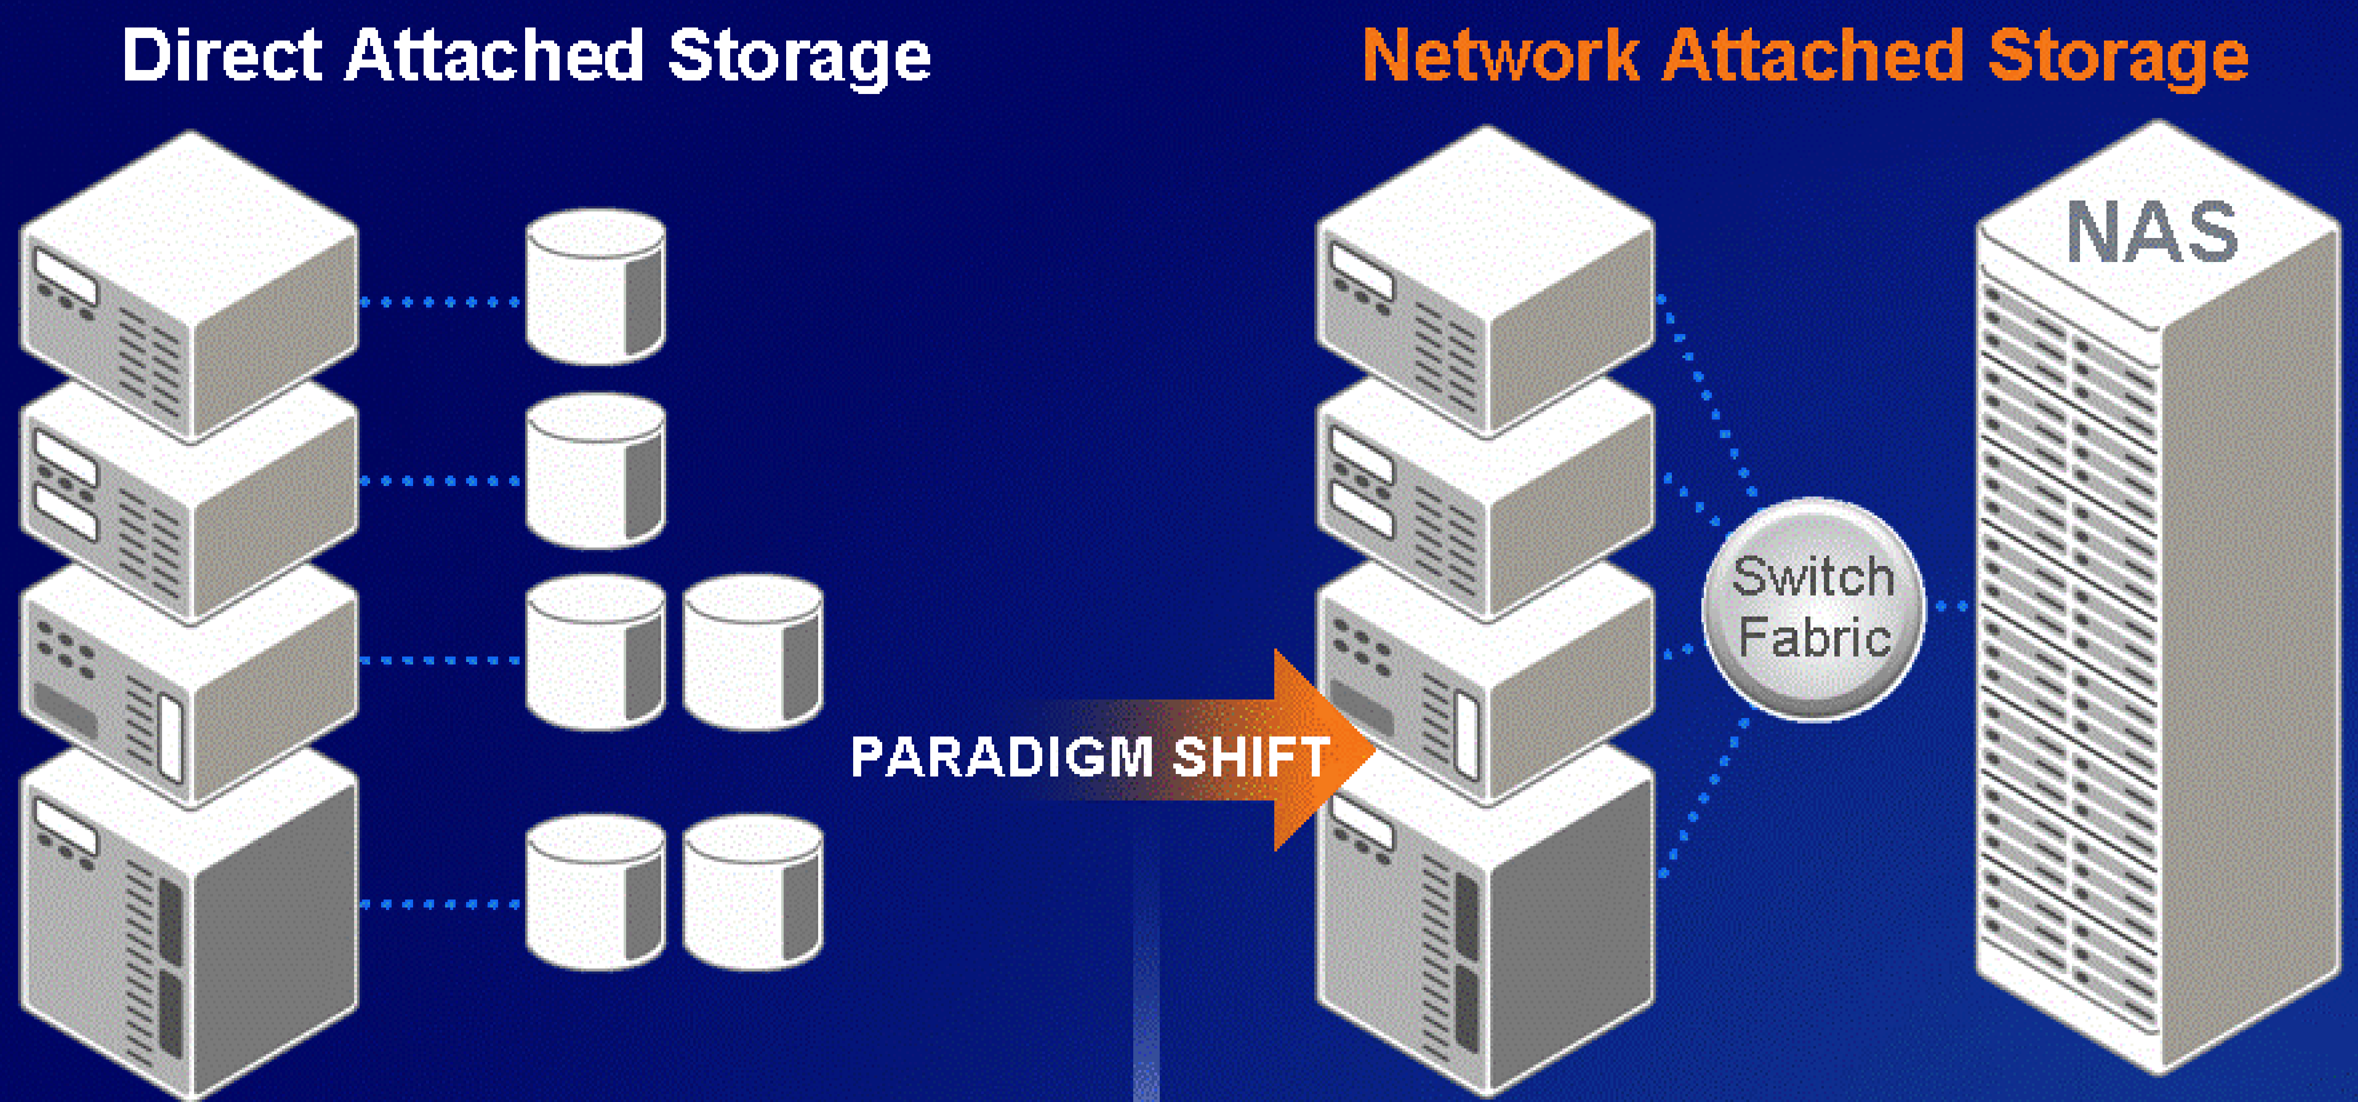

- Partitioned drives
  - 1000 x 1GB drives vs. 1TB
- Management complexity

- Transparent separation
- Concentrated, scalable capacity
- Heterogeneous support
- Significantly reduced TCO

Modern Three-Tier Architecture

There has to be a better way!

The Real Cost of Deployment

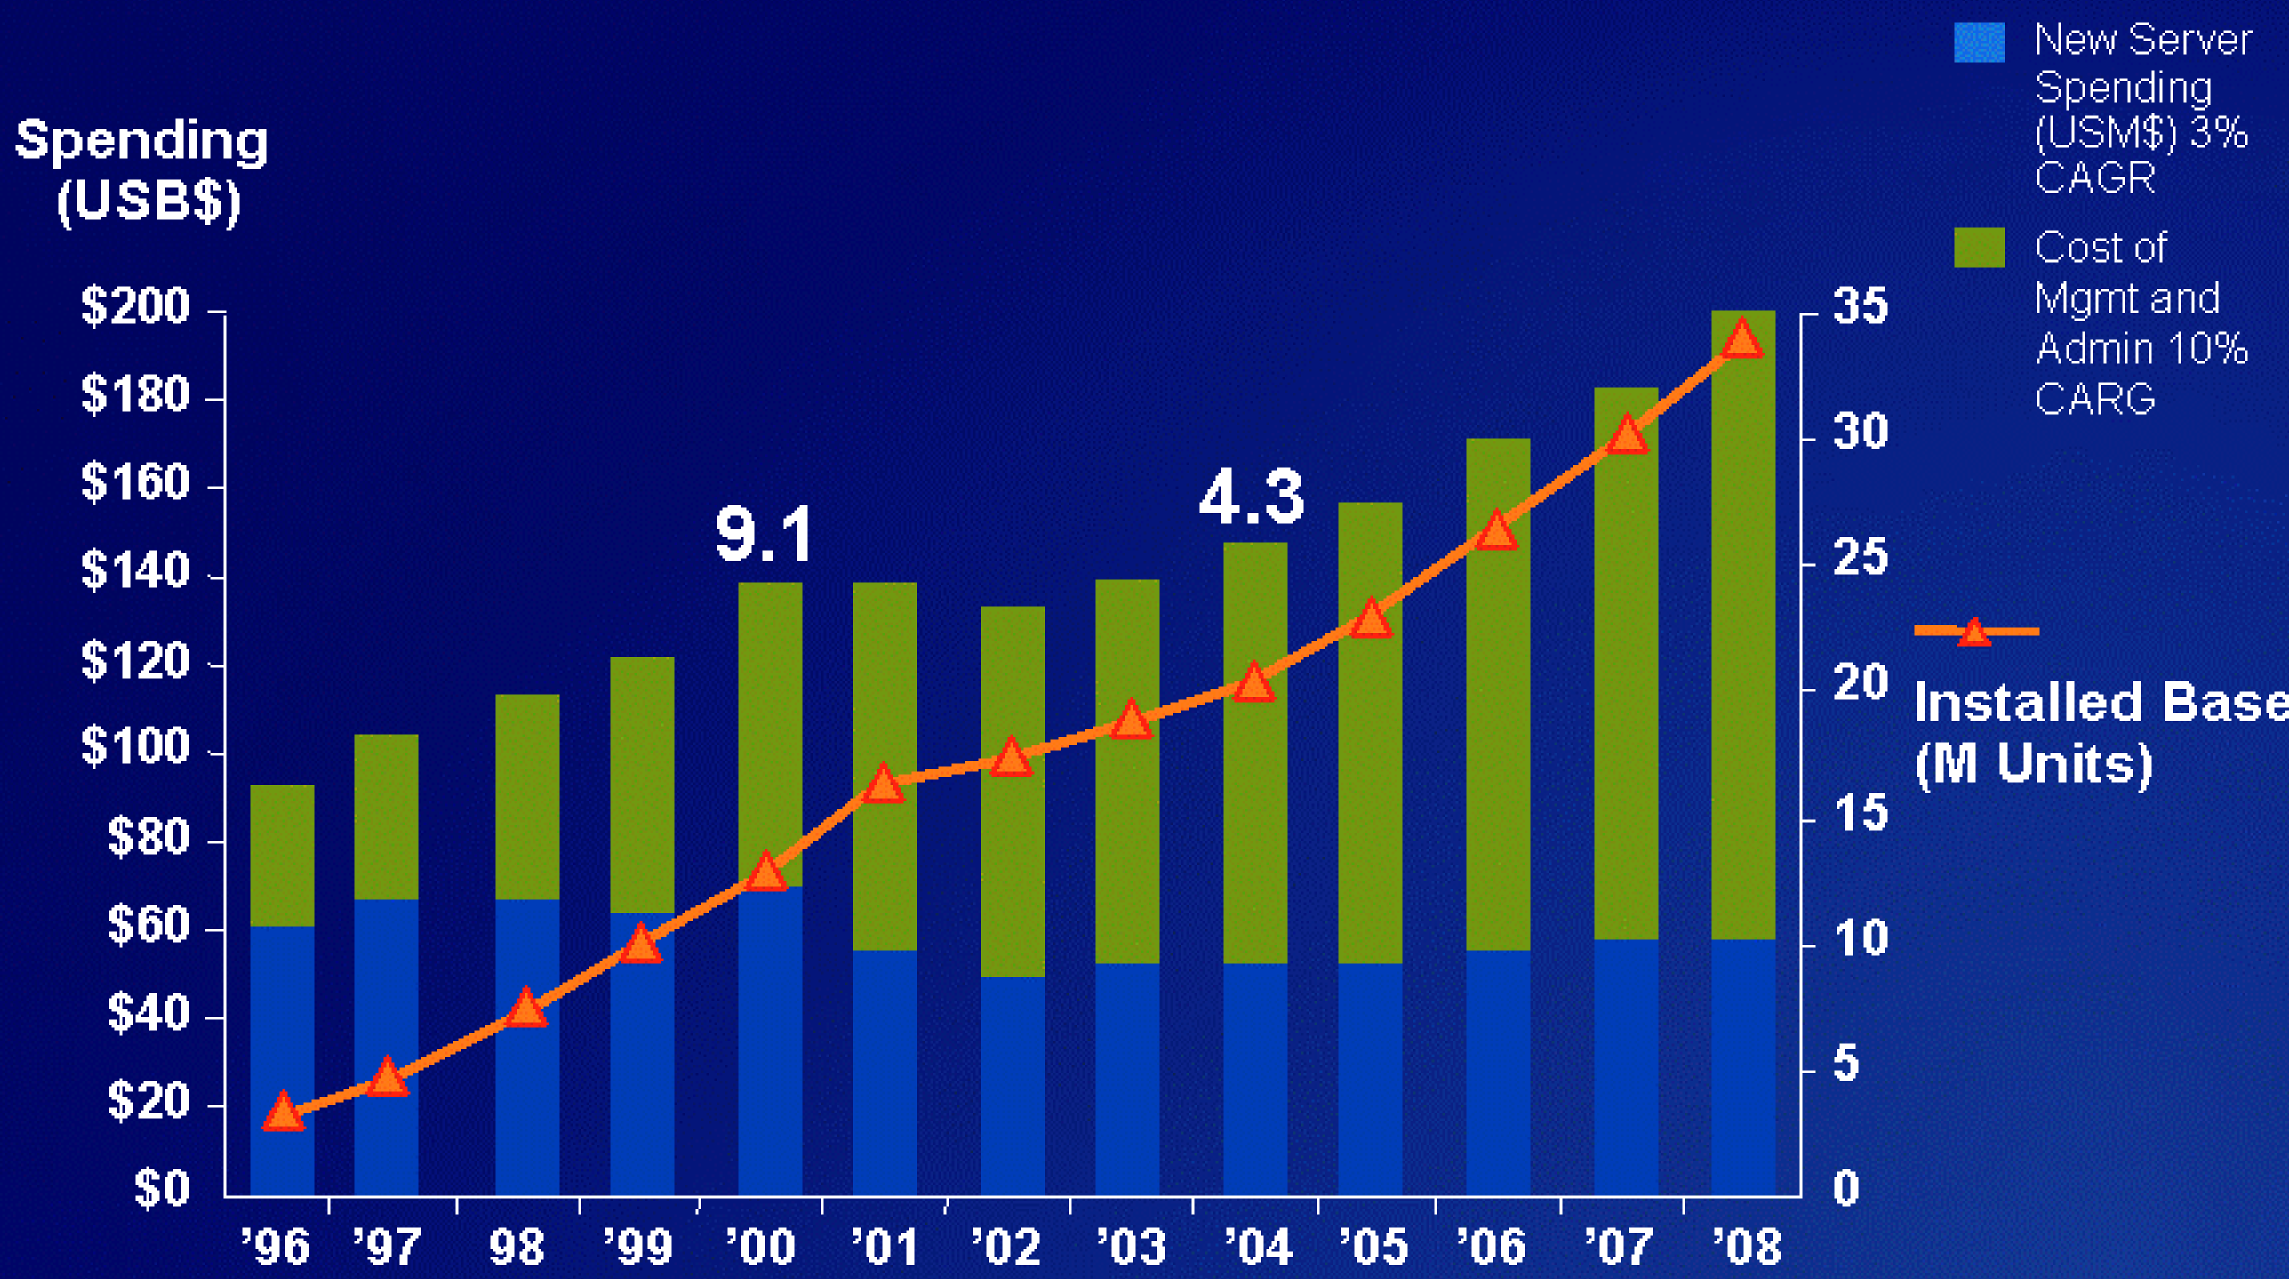

* Source: IDC, July 2004

The Better Way

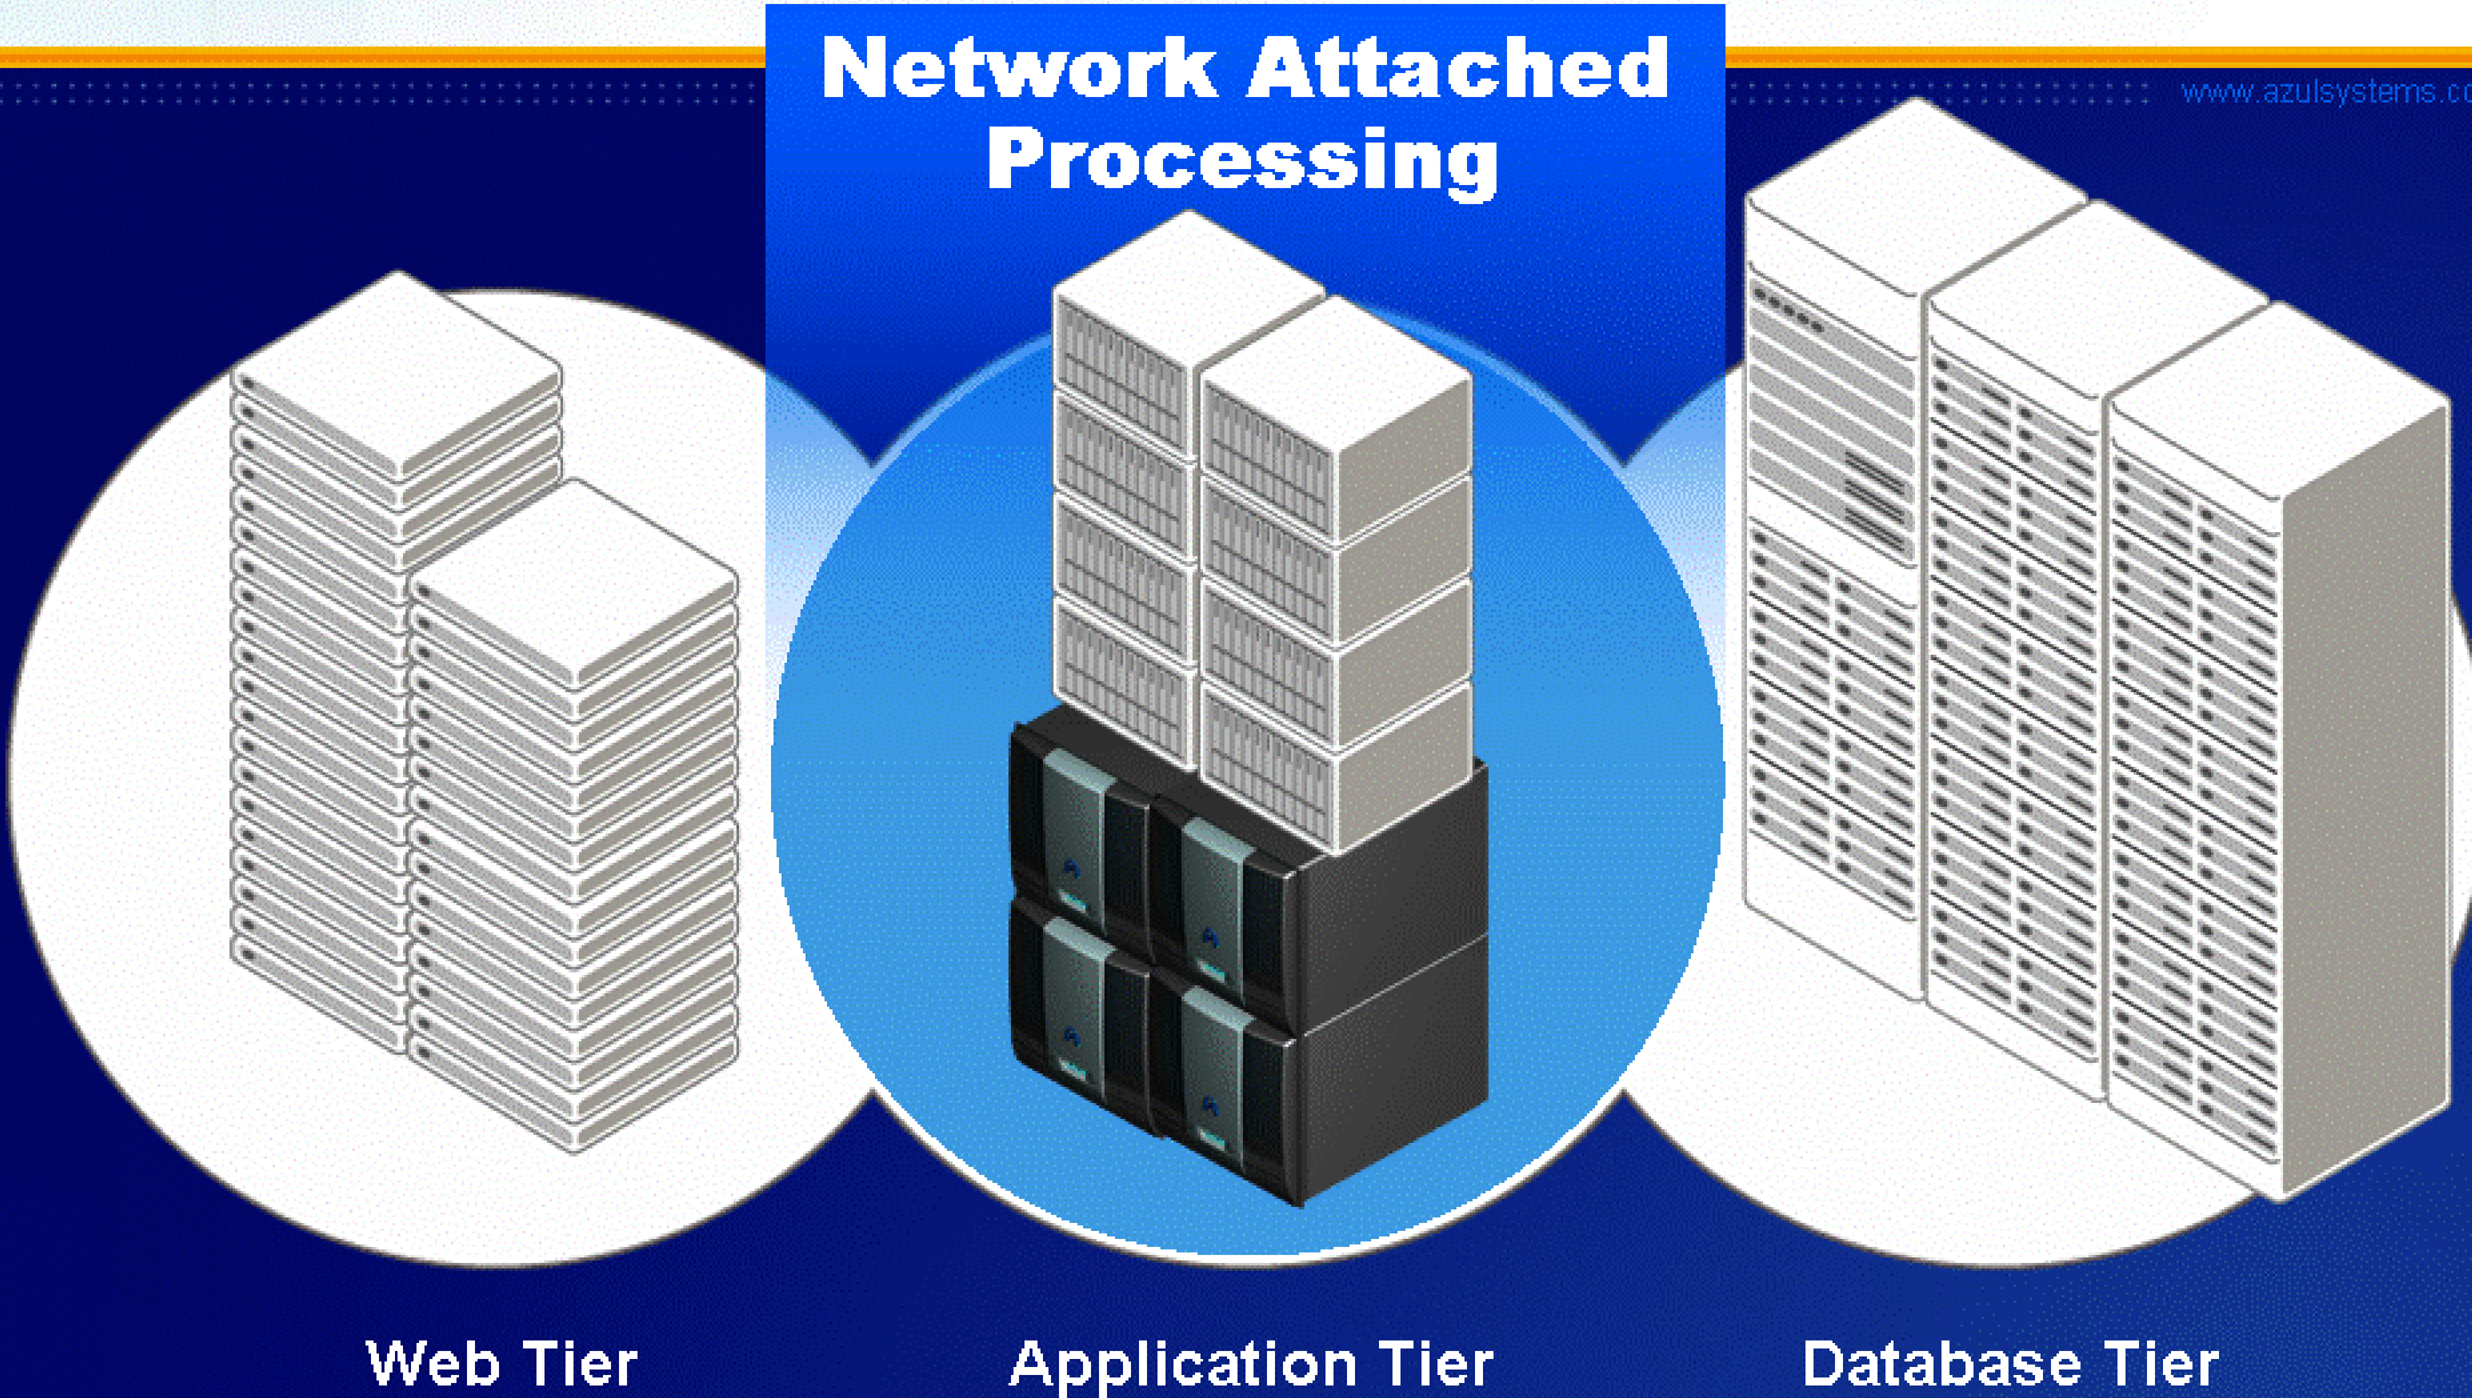

Where Are the Apps?

AZUL SYSTEMS

www.azulsystems.com

Strong Migration to Inherently Multi-Threaded Application Virtual Machines

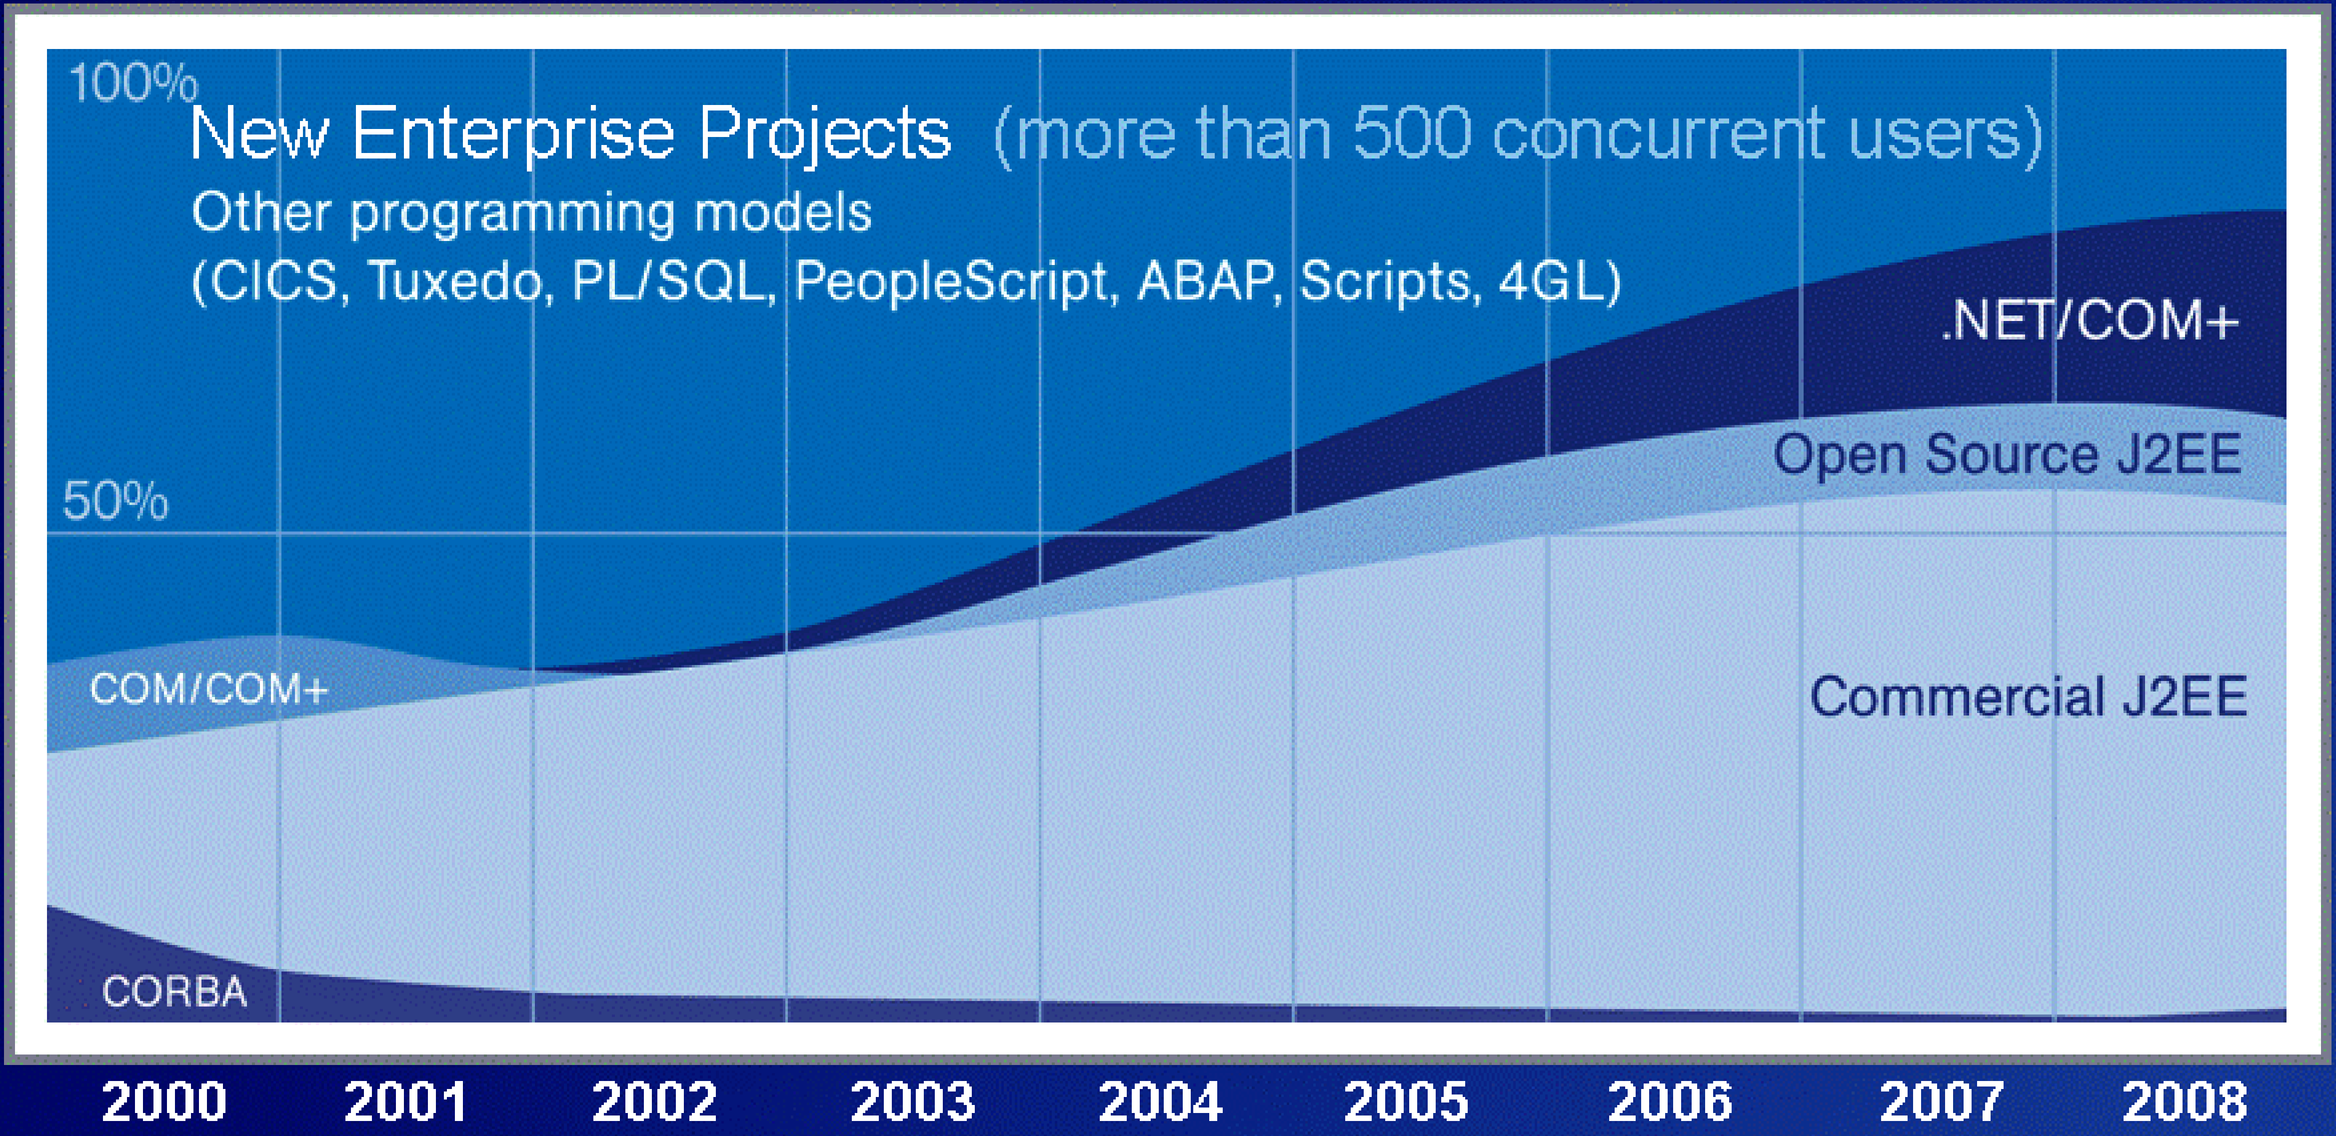

Source: Gartner, September 2004

Azul Vega™ Processor

- Purpose-built application processor
- Instruction set not exposed
  - Cores optimized for VM execution
  - Object-oriented operation assists
  - "Pauseless" garbage collection
  - Optimistic Concurrency
- 24 coherent cores per chip
  - Clusters of 8 sharing a common L2 cache
- Hierarchical caching on-chip
  - L1: 16KB Instr, 16KB Data per processor
  - L2: 1MB shared by cluster of 8 processors
  - Hardware cache coherency
  - Memory consistency optimized for Java
- Glueless support for multi-chip SMP
- Comprehensive RAS features
  - Error detection/correction on all memory elements & links, predictive failure analysis

Azul Compute Appliances

AZUL SYSTEMS

www.azulsystems.com

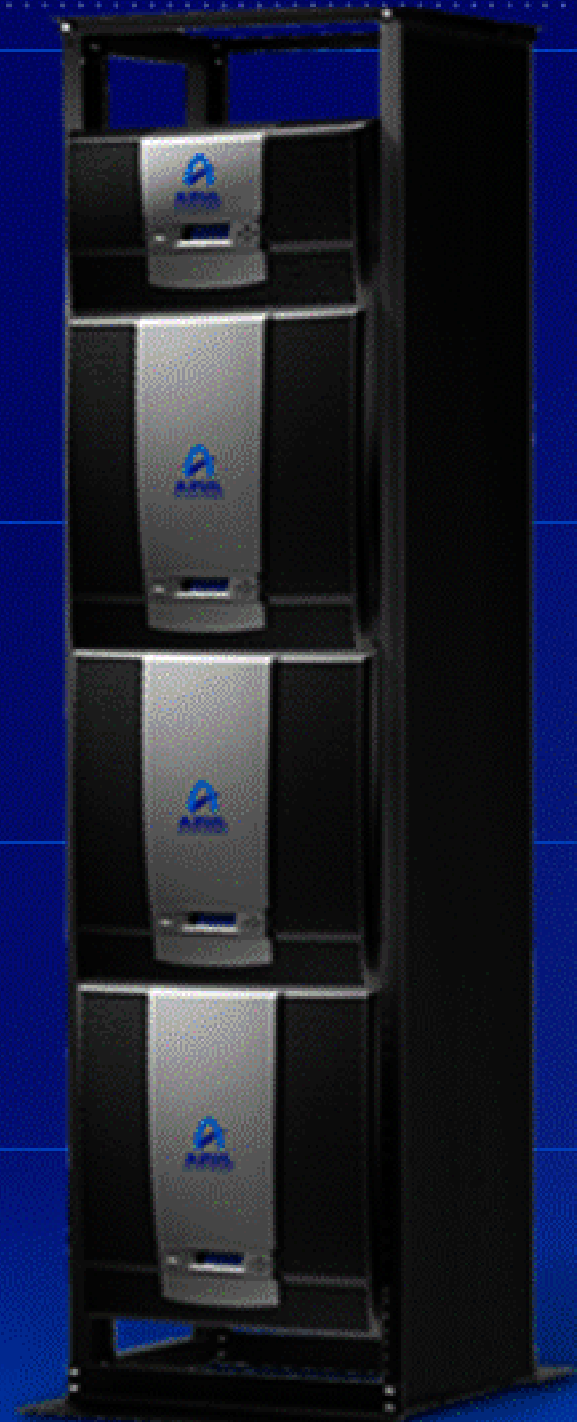

Massive Capacity, Densely Packed

- Up to 384 cores and 256GB memory per appliance
- SMP design with 10ms response to workload spikes
- Each appliance only 11U (19.25”), ~2700 Watts
- 1248 cores per rack—Kilocore deployments!

Removes Classic JVM™ Limitations Around

- Memory size (up to 96GB heap)
- Response time dependency (Pauseless GC)
- Multi-CPU scaling (Optimistic Concurrency)

Transparent Deployment

- No changes to existing Java applications to tap
- Can be tapped by multiple servers simultaneously
- Operating System agnostic

Easy To Manage As A Whole

- No software to install or configure
- Built in Compute Pool Manager
- Simple to deploy and administer

Network Attached Processing

www.azulsystems.com

OS Agnostic, Transparent, Shared, Managed

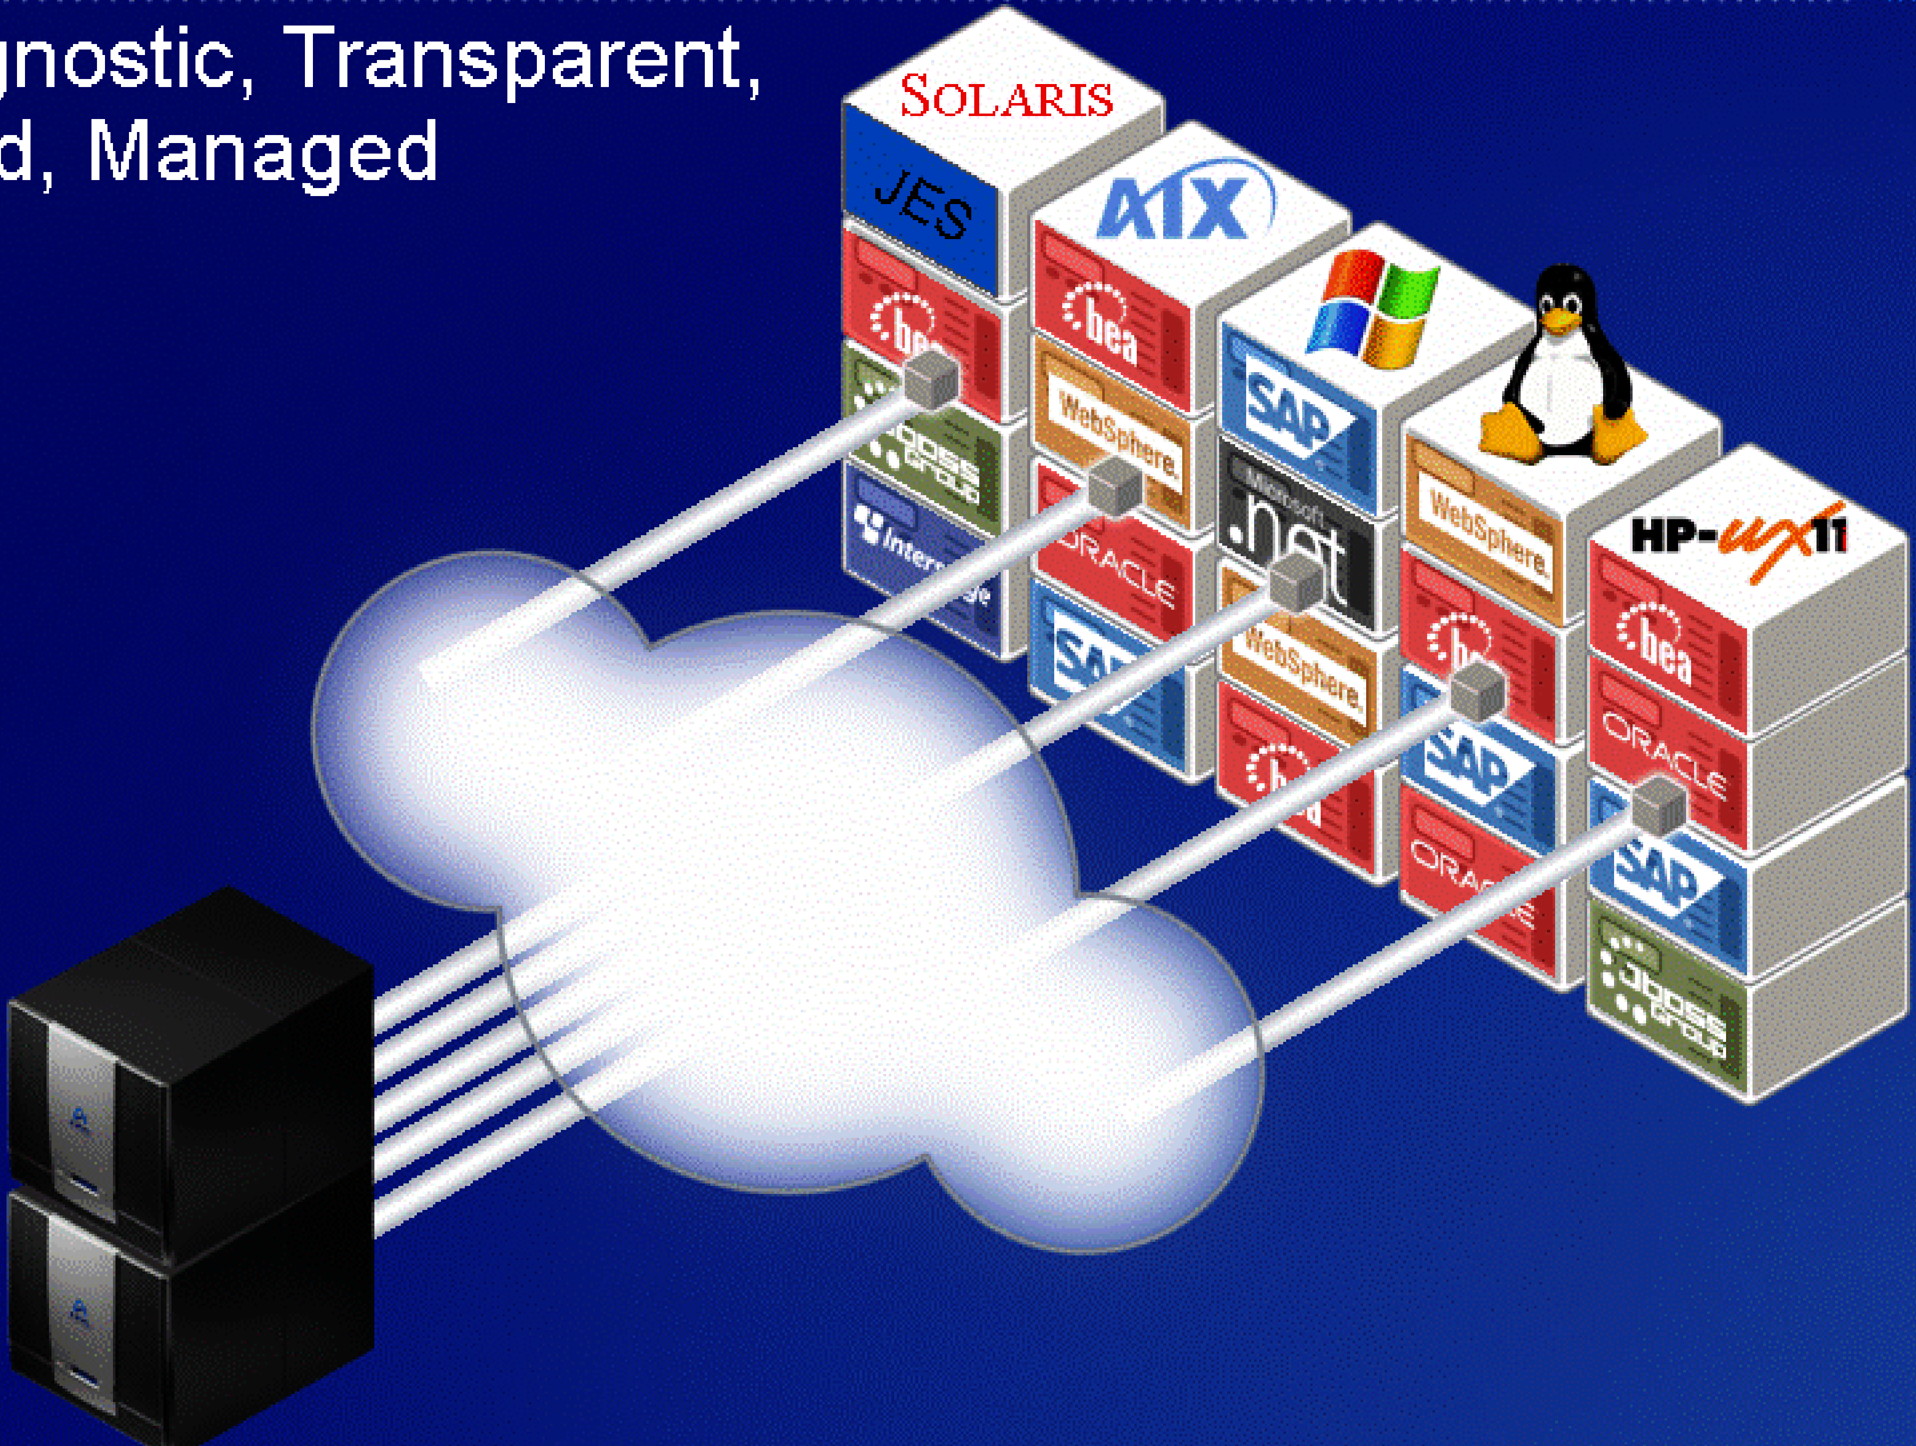

The Value of Shared Capacity

www.azulsystems.com

Traditional Deployment 5 Servers, 5 Separate Applications

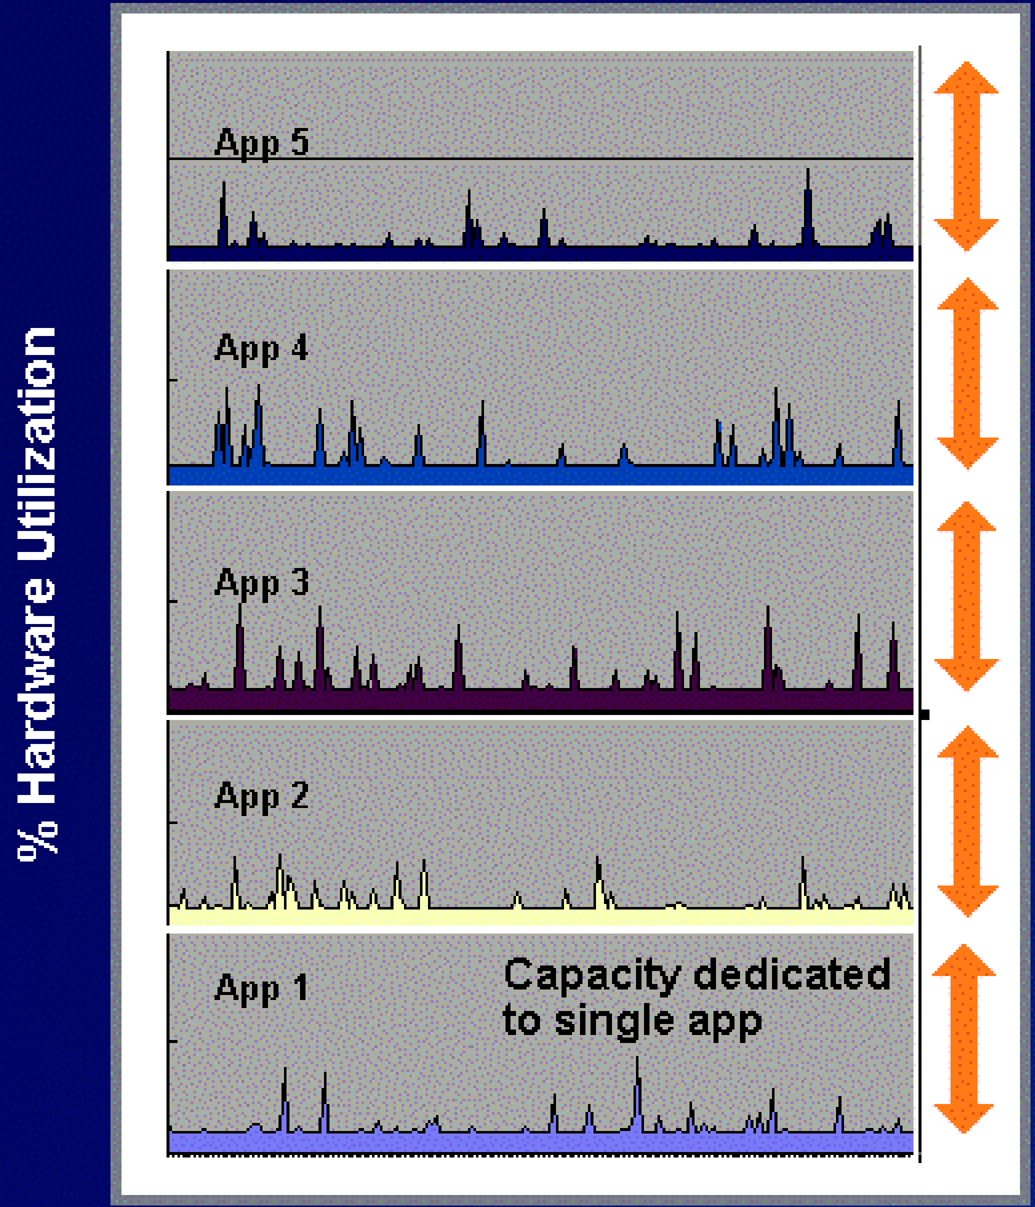

Time | 5 Apps—Utilization ~ 10%

Pooling and Sharing of Capacity

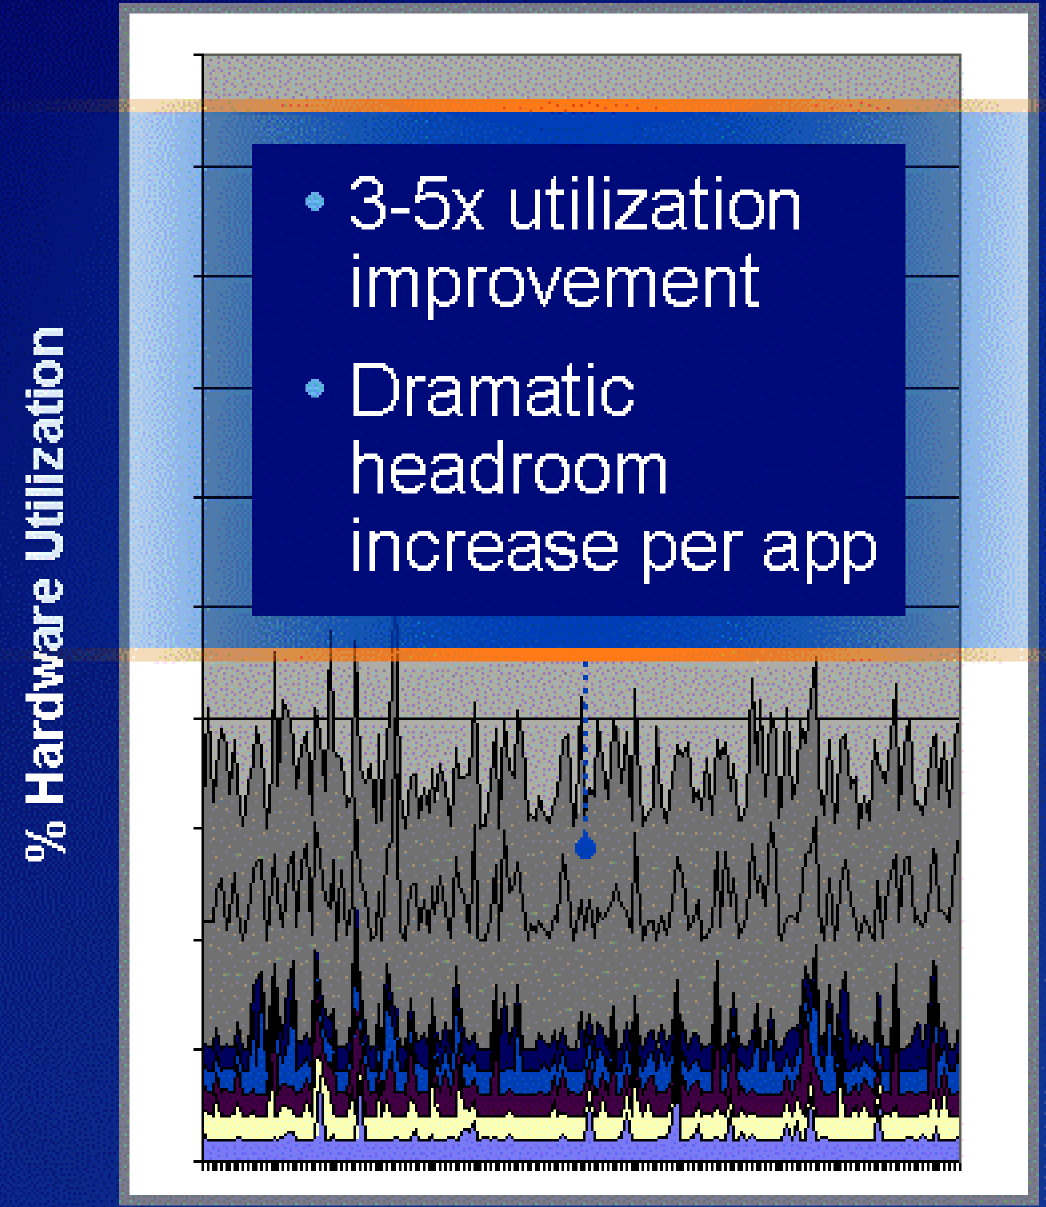

Time | 15 Apps—Utilization ~ 30%

www.azulsystems.com

Software Licensing

AZUL SYSTEMS

Did you say 384 Cores?

How do you deal with software licensing?

- Azul solution to ISV SW licensing
- License the server, not the pool
- Fits a shared-resource model
- Correlates with usage
- Supports mixed work loads
- Lets customers grow into the pool
- Compatible with movement towards Utility Computing

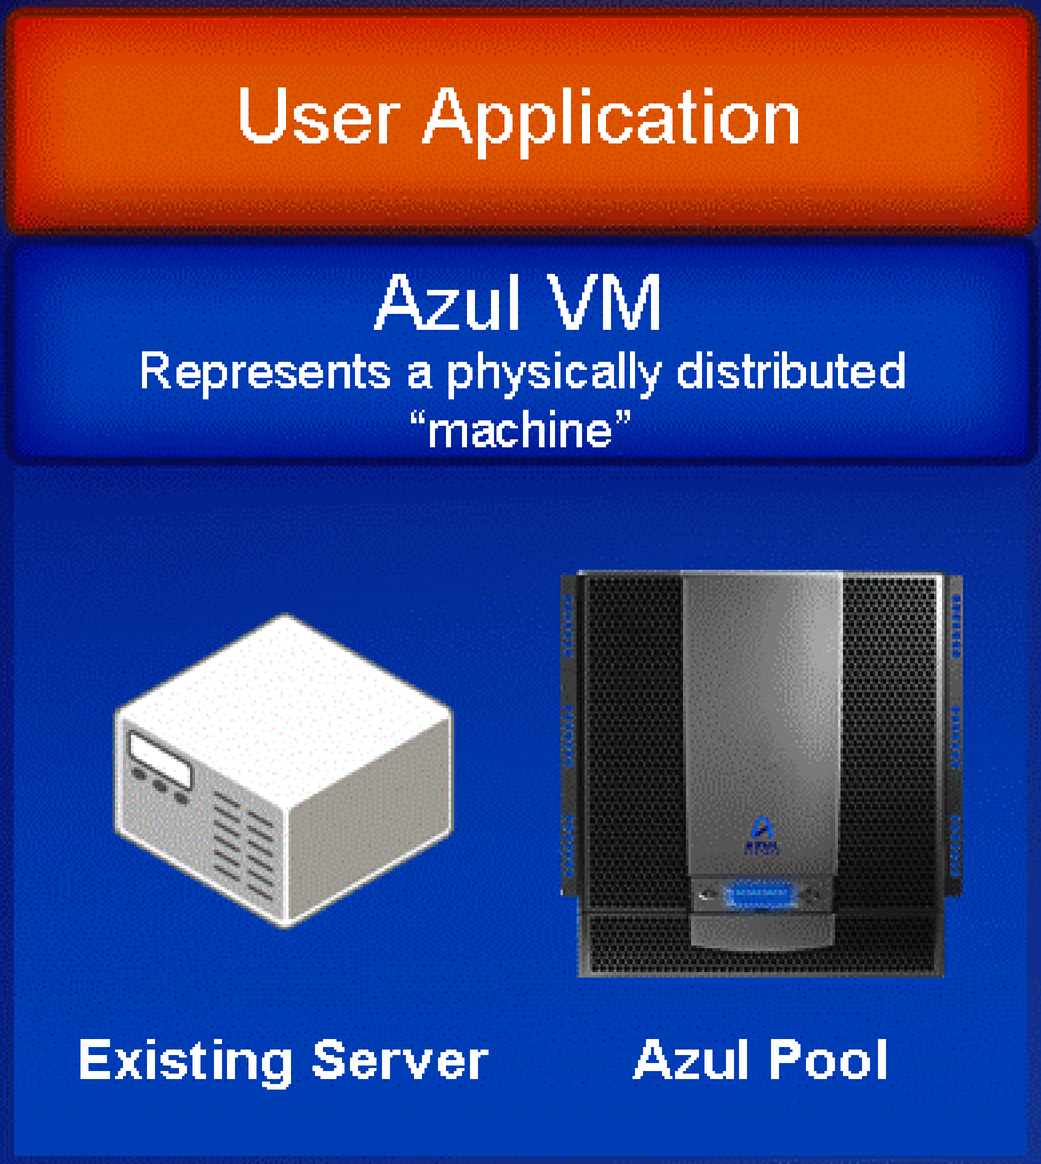

Software Licensing on Azul

BEA WebLogic Agreement: 3x Uplift

www.azulsystems.com

3x host CPU license, no BEA license charge for Azul pool

Cost neutral if number of servers is reduced by 3x

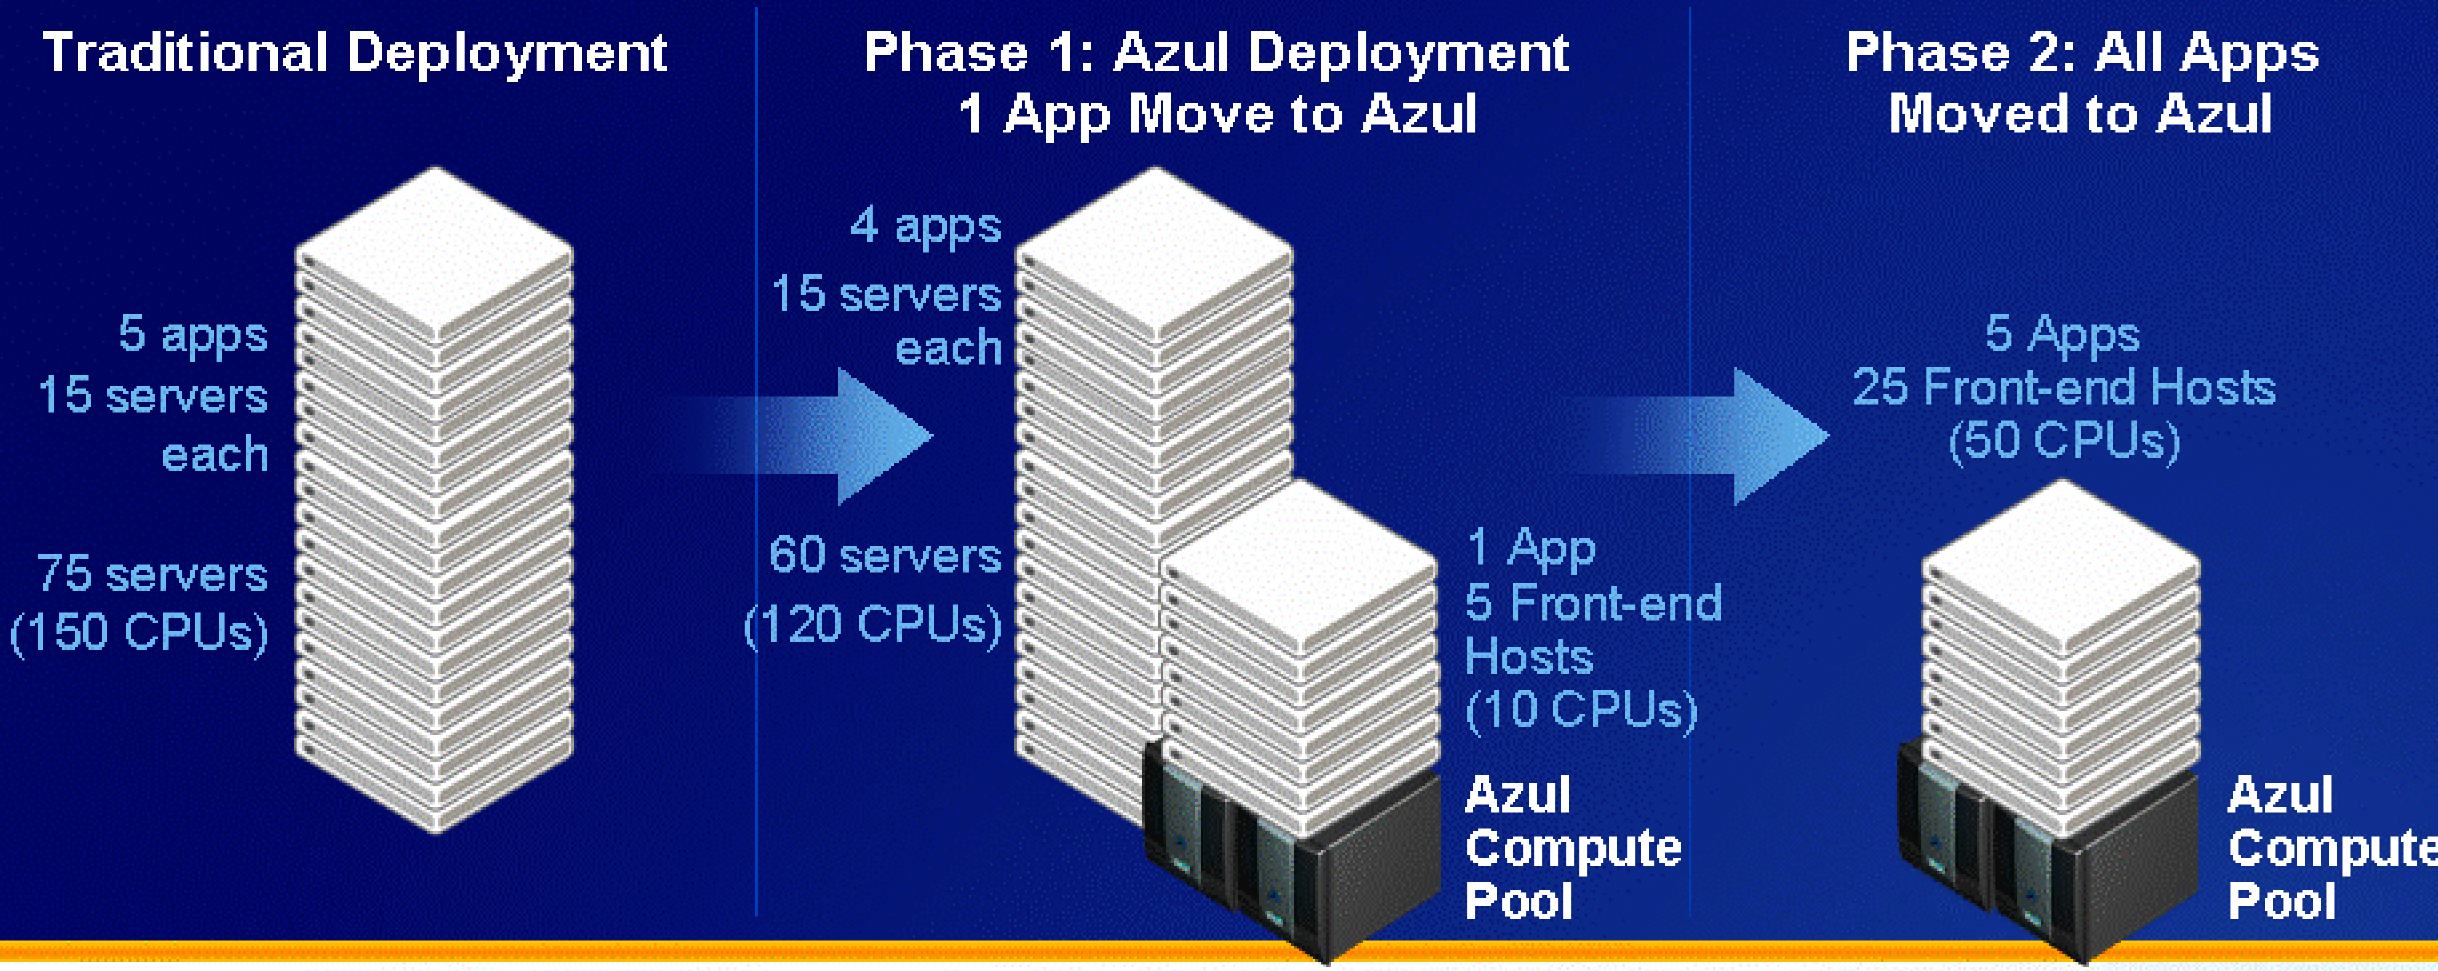

Traditional Deployment	Phase 1: Azul Deployment	Phase 2: All Apps Moved to Azul
150 Licenses **$N/License** **$150*N License Fees**	**13 Licenses** **120 @ $N/License** **10 @ 3*$N/License** **$150*N License Fees**	**50 Licenses** **50 @ 3*$N/License** **$150*N License Fees**

By 2010…

- No "One size fits all" computing architecture
- Processor architectures become bifurcated
  - What's right for the client is not right for the enterprise
  - Client CPUs continue to primarily focus on minimizing thread latency
  - New class of server microprocessors architected to exploit inherent Thread Level Parallelism
- Emergence of throughput computing architectures for enterprise applications
- Prediction
  - The combination of both transistor density and clock frequency increases due to silicon technology scaling will result in the performance of chip multi-processors for multi-threaded applications to double every 18 months

Q&A

UNBOUND COMPUTE™

Thank you

sellers@azulsystems.com

AZUL SYSTEMS

Trademark Usage

www.azulsystems.com

•Azul Systems, Azul, and the Azul arch logo are registered trademarks of Azul Systems, Inc. in the United States and other countries.
•BEA and WebLogic are registered trademarks of BEA Systems, Inc.
•IBM and WebSphere are trademarks of International Business Machines Corporation in the United States and other countries.
•JBoss is a registered trademark and servicemark of JBoss, Inc.
•Microsoft and .NET is either a registered trademark or trademark of Microsoft Corporation in the United States and other countries.
•Oracle is a registered trademark of Oracle Corporation.
•Sun, Sun Microsystems, Solaris, Java and all Java based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.
•SAPand Netweaver are trademarks or registered trademarks of SAP, Inc. in the United States and other countries.
•Other marks are the property of their respective owners and are used here only for identification purposes.

